Fermilab

TM-1429
2280.000

Report of the Census Task Force on Beamline Control System Requirements

E.J. Barsotti, J.F. Bartlett, V.D. Bogert, F.O. Borcherding,
J. Butler, P.C. Czarapata, W.J. Spalding, and A.D. Thomas

Fermi National Accelerator Laboratory
P.O. Box 500
Batavia, IL 60510

October 2, 1986

Fermilab

REPORT OF THE CENSUS TASK FORCE ON BEAMLINE CONTROL SYSTEM REQUIREMENTS

E. J. Barsotti, J. F. Bartlett, V. D. Bogert, F. O. Borcherding,
J. Butler, P. C. Czarapata, W. J. Spalding and A. D. Thomas

October 2, 1986

TABLE OF CONTENTS

1. INTRODUCTION

In January of 1986, the Head of Research Division appointed a special Task Force to study the experience with the present beamline control system and to make recommendations for future directions in this area. The charge of the Division Head and the list of Task Force members are included in Appendix I.

In order to carry out its assignment, the committee conducted a series of meetings in which it discussed the controls situation in general and the best way to approach the user community. The various groups of users were identified and a letter was written to representatives of these groups asking questions concerning the patterns of use of the present system, deficiencies that should be corrected, projections of future needs, etc. Since we recognized that not all of the control systems users had thought long or deeply about controls issues, we designed the letters to be provocative and included questions in each that we hoped would encourage the groups to think well into the future. These letters are included in Appendix II of this report. The committee then scheduled meetings with each group to discuss the response to these questions and freely invited discussion that went outside or beyond the content of our letters. A list of the meetings is also given in Appendix II. Two groups, the Operations Department and the experimenters required two meetings to fully explore all issues. At the end of each meeting, a member of the committee summarized, in writing, the discussion. These summaries are included in Appendix III. In many cases, written material was submitted as part of the proceedings. This material, without any editing, forms Appendix IV of this report.

After all meetings were completed, the committee then reviewed the material which had been gathered and defined certain general topics on which it felt conclusions ought to be drawn. These conclusions are in each case a synthesis of requirements of essentially diverse user communities who interact with different parts of the control system and use that system in fundamentally different ways. In drawing conclusions, we have attempted to identify which user community is the driving force behind the requirements in each area.

Since this document is intended to set the future direction for the beamline control system, the committee felt that it was necessary in some cases to look beyond the specific requirements and needs expressed by our present users and to incorporate some of our perceptions about the evolution of user needs. We have attempted to clearly identify in this document any recommendations that are being made by the committee based on its own insights rather than on information obtained from the users.

2. CONCLUSIONS

2.1 Current Problems

The present beamline control system is not adequate to support the Tevatron Fixed Target program. It does not support an adequate number of users nor does it provide them with services required to efficiently tune and monitor the performance of the beamlines. This results in lost beam time whose value far exceeds the cost of remedying these problems.

2.2 Time scale and laboratory-wide view

Many of the above mentioned difficulties are the subject of work already being undertaken by the Research Division. This work is regarded as remedial in nature in this report. It is assumed that the time scale for the remedial work is short; perhaps less than two years. The general scope of the remedial work is assumed to be to provide services roughly equivalent to those of the Accelerator Division Control System. A five year plan with the goal of achieving a system comparable to that used to operate the accelerator today is not acceptable; it would be five years out of date when completed. A five year plan based upon the user needs outlined in this document should seek to avoid repetition of work which might be accomplished in parallel in the Accelerator Division. The Laboratory and various Division management groups have a distinct responsibility, in our opinion, to foster such uniformity of endeavor.

2.3 Systematic Upgrades And Specific Recommendations

The Research Division should undertake a systematic program of upgrading its beamline control system. The effort should include upgrades to hardware -- power supply controllers, beam detection and monitoring hardware, spill monitoring hardware -- as well as greater computing capacity and better software. During these upgrades, the implementors should maintain awareness of similar activities, past or present, elsewhere in the laboratory in order to minimize duplicated effort wherever at all reasonable. Recommendations in specific areas follow.

2.3.1 Networking

The beamline control system should support all devices required to control and monitor the beamlines and their associated utilities. The various components of the system should be networked so that information from any component is available anywhere in the system.

2.3.2 User Interface

In order to support a diverse community of users, including experimenters, beamline and cryogenic operators, beamline physicists and technicians, the control system must have a vastly improved user interface which includes color graphics with a pointing device like a trackball or mouse. At least some terminals in the system must be of the "intelligent" variety, so that they can be programmed to support the activities of particular end user communities. The system should support at least 80 users. Hardcopy capability should be upgraded and made generally available so that documentation of beamline "tunes" and running conditions can be improved.

2.3.3 Monitoring Requirements

The beamline control system should monitor essentially all important beam line devices on at least a per spill (or per ping) basis. This could amount to 2,000 - 3,000 devices. The monitoring should include checking beam positions, shape, and intensity.

2.3.4 Application Languages

A powerful application programming environment must be available. For the sake of efficiency, it should be based on compiled languages; FORTRAN and C must be provided. BASIC should be provided if an acceptably efficient implementation can be found.

2.3.5 Data Logging

The system should provide facilities for long-term and short-term data logging and for "downtime" logging.

2.3.6 Internetworking

The system must be able to exchange information with experimental data acquisition and control computers as well as the accelerator control system, ACNET. This includes the ability to read devices from those computers as well as restricted abilities to set some devices from them. This interconnection should use commercially available network hardware and software protocols if at all possible.

2.3.7 Data Acquisition For Experiments

The beamline control system should be able to act as an extension of an Experiment's apparatus by supporting the transfer of data from the beamline to the Experiment. The beamline system should be able to support large amounts of data transfer and should allow a straight forward interface with diverse experimenter systems.

2.3.8 System Security

The only truly secure system is one which does not allow remote accesses of any kind; however, imposing this restriction would severely compromise the usefulness of the control system. Access must be allowed from the experimenters' computers; and, since these are connected via multiple paths to the outside world, remote access through these computers to the control system is always possible. The administration must impress upon the experimenters the need for reasonable security measures on their data acquisition computers, not only to protect the integrity of the beamlines but their experiments as well.

3. DISCUSSION OF CONCLUSIONS

3.1 Scope Of The System

There was agreement among task force members and implicit agreement among the users that the beamline control system ought to include all systems, services, and functions required to operate the beam efficiently. All these systems should be integrated so that data from any subsystem can be available to users anywhere in the system. While all this may be obvious to the designers of the accelerator control system, it has not been the custom in the Research Division. There are several separate systems -- for example beamline control, cryogenics control and radiation monitoring which developed separately and without any plan to connect them together. The need for integration is highlighted by the plight of operators who must interact with all the pieces and must move between various terminals and monitors to do so. In fact, the three areas, Neutrino, Meson and Proton are not even connected so that a single operator has difficulty in handling problems which occur at roughly the same time in two areas. The existence of a separate cryo control system not well supported at the Operation Center has prevented the normal operations staff from assuming responsibilities in this area and has probably forced the lab to use more operators than really necessary.

The implications of this principle for future directions is that hardware and software choices should produce a flexible and expandable system so that functions we do not now know about can be added easily. This also means that there should be ways of adding additional capacity to support new systems.

The collection of existing systems which should be included in the control system are:

- Usual control of beamline devices through an upgraded hardware system with expanded support for new modules;

- Support for Cryogenics control based on the already installed Multibus crates;

- Support for monitoring and control of vacuum devices based on the Accelerator Division CIA crate;

- Radiation monitor and security system status;

- Beam profile devices, especially SWIC data from presently installed SWIC scanners;

- Remote control and monitoring for utilities such as LCW subsystems as well as monitoring of environmental status.

This list of systems will obviously grow and plenty of capacity must be available to support this growth.

In addition to the above systems, the beamline control system must support the following functions:

- A much more powerful and versatile user interface;
- Powerful plotting, graphing and drawing facilities;
- Good connectivity to the experiment data acquisition computers and the accelerator control system;
- Data logging facilities;
- Fast time plotting facilities and possibly transient recording capability;
- Downtime logging;
- Various services to aid in administrative support of the operators such as filing systems for procedures, electronic log book, etc.
- The system should provide a mechanism for preventing simultaneous setting access to a device by more than one user at the same time, possibly by simply notifying users that a conflict has been detected.

Even this list, which is very ambitious when compared to our present system, should not be considered complete. If there is any lesson to be learned from recent experience, it is that *we must be able to handle demands for entirely new services and additional capacity!*

3.2 Basic Services

3.2.1 Parameter Display/Control and User Interface

The user interface supported by the current EPICS system is by means of a monochrome terminal supporting text and graphics. In text mode, the screen offers 24 lines by 80 columns. The screen is divided into a scrolling region for user input of about 3 lines and a display area for a parameter page. A keypad allows one to scroll the PAGE lines vertically, and function keys select various overlays which in some sense scroll "sidewise". Most input is through simple typed commands. Only one format of parameter display is available, independent of subsystem or end user.

The terminals are not connected directly to the beamline computers but instead are connected to the CAMAC system and must be polled periodically. This forces all data and requests to be transferred from the lowest levels of the system, the data acquisition front ends, all the way to the application levels and then back down on most user interactions. This results in the inefficient use of all the computers and the CAMAC link in the present system.

Two of the most frequently voiced complaints of beamline physicists and experimenters are that the system supports only 16 terminals for each area so that sometimes it is impossible to get onto the system at all; and when the number of users is large, (approaching this maximum of 16), system response deteriorates, causing characters to be lost, commands to be mistyped, etc. The number of users seems bound to increase as several more beams are being built. The committee has attempted to look into the future to guess the maximum number of users that need to be provided for. A number in the range of 60 to 80, for all three areas combined, appears to be reasonable. Additional (perhaps "view-only") users from the central laboratory building and the Accelerator might require an additional 20 terminals. Therefore, a total of 80 to 100 simultaneous interactive users should be accommodated in any future planning.

On the present system a graphics window can replace the text display window. Graphics is supported only for fairly specific operations--time plots of ADC and scaler values and correlation plots. No support is provided for graphical representation of systems with readings superimposed or "informative displays", for example of devices in alarm. One major complaint is that the terminals are "nonstandard" so separate "beamline system" terminals must be provided even in locations where many other terminals are available. The terminal hardware is not very reliable. The restrictive use of windowing forces one to use two, and sometimes three, screens to perform beam tuning or trouble-shooting operations. Up/down buttons, provided to permit continuous adjustment of device settings, do not work properly.

There is unanimity among the user constituencies for the need for a greatly enhanced user interface for the control system. In order to efficiently tune, monitor and trouble-shoot the beamlines and their associated systems, the control system must present the user with data on many devices in an easily readable form, and allow the user to perform monitoring, plotting, and tuning of several devices concurrently. This requires much improved graphics and a multi-tasking capability with multiple windows or monitors, at least for those tuning stations in the primary locations such as the Operations Center, cryogenic control rooms, and the experiments. The service buildings could be equipped with more modest terminals which would be used principally for trouble-shooting a particular problem and thus don't require extensive multi-tasking capabilities and multiple windows.

The beamline and cryogenic operators, who control thousands of devices and often need to graphically trouble-shoot several devices simultaneously, need either window management on a large high-resolution screen, or several medium-to-high resolution monitors, with extensive use of color text and graphics. These tools, especially high-resolution color graphics, are particularly useful in representing process control systems such as cryogenics or vacuum systems. Experimenters share the same needs, but since they typically control only their own secondary beam, and do fairly limited trouble-shooting of a few devices, they place less emphasis on the need for color. A monochrome plot is fine for one or two devices (say solid and dashed curves), but is impractical for several devices (small dots, dashes, chain-dots, etc.) where the plot becomes very hard to read. Color allows many devices to be plotted together, and allows various classes of devices or attributes to be highlighted on a page of text. The multi-tasking capability allows the user to work with several parameter pages and/or plots. This is best accomplished using a windowing environment (in which independent processes use different portions of the screen for IO) of the kind used in the present generation of workstations, and indeed these devices are probably appropriate for the primary tuning locations. In places which do not justify the expense of full workstations, multi-tasking might be achieved by switching the screen of a less expensive terminal between concurrent tasks. With the ever-falling cost of "low end workstations" and the increasing cost of custom software, it might be cost effective to use workstations in all locations

In order to minimize typing, and to reduce reliance on a command language, the user interface to the applications programs should be menu-driven, using a mouse or track-ball on the higher performance stations. Knowledgeable users should be able to by-pass the menus by typing in their choice. In addition the beamline operators stated a desire for hardware "knobs" which could be "attached" to a device to provide pseudo-real time control, (meaning probably a fraction of a second in read-back rate). In this context, real-time reading and displaying of SWIC data is highly desirable. The ability to monitor SWIC profiles, and to view several of them in mini-plots on the tuning station screen was stated by all to be highly desirable.

All the user groups wish to make much more extensive use of graphics than is possible on the present system. They want to be able to define logical and arithmetic functions of physical devices, and to be able to plot several of these physical devices either on the same axis or side-by-side on a reduced scale. For each such device or "pseudo-device", a data rate of 10 to 20 points per second is satisfactory.

Three of the groups which were interviewed -- the operators, the beamline physicists, and the experimenters -- reported a class of problem in which several users would attempt to simultaneously control the same device or set of devices. This situation is most likely to occur during the startup when several beamlines are being tuned. The consequence is that the beamline device characteristics appear to be inconsistent and setup time is prolonged. Various solutions were discussed, including (1) the ability to reserve exclusive control of a set of devices for a limited time and (2) issuing warnings to users when there is actual or potential interference. The users expressed no strong preference for one of these solutions over the other.

3.2.2 Hardcopy Devices

In the present system, hardcopy devices are not widely available. The Operation Center is supported with devices which print video displays but in general experiments are not. This makes the documentation and analysis of beam tunes very time consuming.

Each tuning station should be equipped with a printer for text and graphics. These printers can be medium resolution and monochrome, but all color tuning stations should have access to high performance color printers in a few central or high-use locations.

3.2.3 Beam Monitoring Device Support

Beam tuning and monitoring relies on three devices: (1) SEMs (secondary emission monitors); (2) beam loss monitors; and (3) SWICs (segmented wire ionization chambers). These are backed up by numerous specialized devices such as scintillation counters and Cerenkov counters. The SEMs and loss monitors can be read by the control system, displayed on parameter pages, and used for beam tuning. SWIC data is *not easily available* to the control system. The data is accumulated by a local microprocessor and stored in its memory. The data is then transmitted through the lab-wide TV system and can be viewed on TV monitors. The control system can read the scanners through an 032 interface but the software support for this activity is limited. One reason for this is due to the heavy burden placed on computer resources in order to support these transfers. Thus, we are in a situation where the beamline control system in general does not know the position or shapes of the beams, and therefore cannot monitor them. Operators, experimenters, and beamline physicists all strongly desire that beam profile information be generally available in the control system. SWIC profile shape parameters should be available to the monitoring system and for page displays. All SWIC profiles should be available through the control system on graphics terminals. Even when SWICs are properly supported in the control system, we still see the need for having some of them available to the lab-wide TV system. However, only the most "popular" SWICs need to be on the TV system. This should permit us to live within the limits of the present TV system (and its presently planned upgrade).

At present, the Research Division has no standard hardware for monitoring beam spill structure. We rely heavily on information provided by experimenters from a variety of devices in their apparatus. Experiments vary widely in their susceptibility to bad spill quality and to their sophistication in measuring and describing it. Since spill quality problems can be generated in the accelerator, switchyard, or experimental areas, it is often very difficult to identify their source. This makes it very important to have a complete and well defined approach to measuring the spill quality. The Research Division should develop a standard spill quality monitor.

The three types of detectors described above do provide the basic information required for tuning and monitoring beams. However, a variety of other detectors might also be useful. Examples are RF intensity and position monitors, split plate SEMs and loss monitors, and systems of proportional chambers which would allow systematic, correlated studies of beam phase space. Some of these devices exist but need to be supported within the control system, some have been developed by the Accelerator Division but aren't used in the Research Division, and some would need to be developed from scratch. In all cases, we should develop devices which make it easier to trouble-shoot problems, such as beam motion, spill structure, or beam phase space and transmission problems. (Some anticipated needs for particular beams or experiments are described in Section 3.5.3 below.)

3.2.4 Monitoring

The comments on alarms and monitoring point to the following requirements for the future:

- The monitoring system should have access to information about beam position profiles, intensities and spill quality.

- There is a consensus among the control system users that there should be a monitoring facility which monitors all important devices each accelerator cycle with standard limits, and reports alarms to the beamline and/or cryo operators. Typical loads would be about 100 devices per beam including power supply readbacks, analog and digital status, scaler and loss monitor readings. Beam profile information, for example, means and standard deviations, would add an additional load. One can estimate for 15 beam lines a load of 2000 to 3000 readings per beam cycle.

- The experimenters and beamline physicists need the ability to set their own limits on devices, and to have the alarms addressed to their location.

- It is essential that the alarm messages be formatted clearly and concisely with all devices in alarm listed on alarm pages, according to their location or subsystem. An audio and video message sent to the user's terminal would indicate to the user that the device has entered its alarm state. The alarm displays should be accessible on any terminal, and permanently displayed in some locations as "comfort displays".

- It is desirable that the alarm system should be able to cause designated user-written application programs to run in response to an alarm. Delays of seconds or even minutes might be expected in activating such a program. Also, only a relatively small number would be permitted. *The intent is not to attempt continuous closed loop control with this type of facility.*

- The alarm system should be able to make information available to application programs which could then handle special displays employing color to highlight problems.

- There should be a standard way to interface attention-getting devices such as sonalerts to the alarms system.

- There should be optional logging of alarm messages to disk files and it should be possible to print out the message file.

- The alarm system should support the monitoring of calculated or derived quantities such as scaler ratios.

- It is desirable to have conditional alarms, that is, alarms which are only declared if other devices also satisfy certain conditions (for example alarm if this scaler ratio is less than x provided the intensity of the machine is greater than y).

- The alarm system should have a scheduler so that devices need be checked only as often as is sensible.

- It is desirable that alarms should be prioritized and grouped so that if many problems strike at once, the messages are as meaningful as possible.

3.2.5 Save/Compare/Restore

The ability to save, compare, and restore device readings and settings enables operators, beamline physicists and experimenters to benchmark good beam tunes. This function is carried out reasonably well by the present system as far as beamline power supply and device settings go. While no "compare" program is available, they can (and have) been written as CBASIC applications for some beamlines. The main deficiencies in this area are the lack of a *convenient* means to benchmark the actual operation of the beam -- including making hardcopy records of SWIC profiles (currently, this is done by the cumbersome method of taking pictures of the TV screen SWIC displays with a Polaroid), intensities, and loss patterns. In the future, we need a system at least comparable to the accelerator's "SWIC TREK" application program for displaying and recording beam profiles and a good way of collecting and recording intensity and loss profile data along with beam settings. The latter could be accomplished by special application programs. This is another area in which a significant improvement in facilities for making hardcopy of displays is urgently required.

3.3 Application Programming Environment

The designers of the EPICs system understood from the beginning of the project that they had to provide an environment in which operators, beamline physicists, and experimenters could develop and run applications that were unique to a particular system, beamline, or experiment. Their response was to make an *interpretive* programming language, CBASIC, the fundamental user interface both for interactive control and for running of applications programs. The developers of the Accelerator system took a fundamentally different approach. They provided the services required to develop programs in higher level *compiled* languages on a VAX and to run them on specially equipped dedicated PDP-11's called "console computers". In effect, they support dedicated workstations for applications.

The main writers of applications under EPICs have been operators, beamline physicists and experimenters. On the positive side, CBASIC has been easy to learn and offers an integrated editing, programming and debugging environment. On the negative side, the documentation is viewed as inadequate; the implementation of the language itself lacks many rudimentary features commonly found in BASIC on personal computers, there is no easy interface to graphics, and the availability of CPU cycles for BASIC programs is poor. The inefficiency inherent in the use of an interpretive language has produced frustration both for users, who suffer from poor response, and from the maintainers of the system who see the applications produce degradation in fundamental control system services such as terminal response, parameter page update, etc.

Given the inefficiency of an interpretive applications language, the Task Force attempted to determine whether a compiled language would gain acceptance among the users. The operators listed the availability of a compiled BASIC and other compiled languages such as C as "essential" but stated that an *interpretive* version of BASIC might be "useful". Beamline physicists and experimenters said that they do not feel that anyone would be discouraged or intimidated by the use of a compiled language for applications. Based on these responses, the Task Force concludes that future systems should provide a powerful applications programming environment based on compiled languages. The implementation of an interpretive BASIC in future systems is considered desirable but the decision on whether to implement it, and with what priority, should be made by the system developers based on an analysis of its benefits as weighed against the effort to produce a BASIC that remedies that objections to the present CBASIC.

3.3.1 Programming Languages

Tastes vary among various constituencies. The experimenters program mostly in FORTRAN as do many of the beamline physicists. The operators prefer C and BASIC. It is obvious that all these should be made available to applications programmers if possible.

3.3.2 Programmer Discipline

Programs written to assist in beam tuning, hardware trouble-shooting and analysis, etc. may be as essential to the efficient operation of a beamline as some of the hardware. It is therefore necessary to have a disciplined and organized approach to program development, documentation and maintenance. This sort of discipline has been lacking in the Research Division. In fact, it was difficult even to discuss with the beamline physicists and experimenters systems like the Accelerator Division code capture system because the whole concept was alien to them. The Committee believes it is up to the Research Division management, working with the system developers and seeking advice from the Computing Department and the Accelerator Division, to encourage a disciplined approach to applications programming. Responsibilities should be assigned to individuals or teams, information transfer between users should be encouraged, etc. The essential point here is that the system derives much of its power from user written applications and an organized, disciplined approach must be taken with respect to the creation of these programs, their maintenance, their documentation and their dissemination.

3.4 Data Logging

The present control system has virtually no provision for long-term data logging. Short-term data logging can be accomplished via CBASIC but is restricted and cumbersome. The various groups made requests for long-term, or continuous data logging in the following areas:

3.4.1 Beamline Intensities

Intensities should be logged on a per-pulse or per-ping basis for all primary beams and secondary beams. Recording of "good data" indicators from experiments is desirable. The volume would be of order 100 numbers per beam spill. Data must be kept for a long time (perhaps years) but could be migrated onto tape very frequently.

3.4.2 Downtime Log

Research Division management emphasized the importance of a computerized downtime log. This would be used to assess subsystem reliability to help eliminate major sources of beam downtime in a systematic way. Used in conjunction with the Intensity Log, it would help Program Planning get a good picture of what was going on. It was generally agreed that the Downtime Log and an Intensity Log could be separate control system functions and could be correlated "by eye" where necessary. There was also agreement that a "manual entry" system would be acceptable even though the operators might occasionally forget to enter something or might enter just the most "upstream" problem when several occurred at once. The system should have facilities to compute the downtime for each beam by major system (e.g. power supplies, Cryo, interlocks) and should in general have good report generation facility.

3.4.3 Cryogenic Systems

These systems require long-term data logging of a great number of devices in order to spot long-term trends, evaluate system performance and to tune for optimum efficiency. The total number of devices that could be logged numbers in the thousands. However, after some discussion, it was agreed that the minimum requirement was 50 items per refrigerator installation for 5 to 10 installations recorded every 15 seconds. They emphasized the importance of being able to "play back" the logged data on the control system itself without inhibiting data-logging and to have user-friendly tools within the control system for studying the data and modifying the list of items to be logged.

3.4.4 Radiation Safety

Radiation Safety would like to log on the order 15 SEMs, 30 paddle counters, and about 140 chipmunks every spill to a circular buffer that would turn over (wrap around) about once per hour. This would be useful in case of an incident.

3.4.5 Site Operations Department

This group is interested in logging "most" of their devices at a frequency of once per shift or once per day.

3.4.6 Short-term Logging Requirements

Several groups requested the ability to do short-term logging. This means the ability to log for several hours, and perhaps several days, a few devices to try to investigate a particular problem. The following requests were made:

- Beamline physicists have requested the ability to log several devices for a "few" shifts to trouble-shoot long-term drifts, intermittent problems, etc.

- Radiation Safety would like to log not more than 20 devices for periods of 2 to 3 days on a per pulse basis for studying new beams, new operating modes, etc.

- Site Operations also expressed an interest in short-term logging for trouble-shooting a few devices.

3.4.7 Transient Recording

Some groups, in particular Cryo and Electrical, have a need for recording rare or intermittent events. Capturing the sequence of events during a quench is an example of such a need. This kind of problem can be handled best by a smart A/D module or specially designed transient recorder supported by software for presentation and evaluation of the results. The new system should accommodate this need.

3.5 Data Interchange With Experiments

3.5.1 Requirements and Current Problems

It has long been recognized that it is necessary to include the ability for experiment data acquisition computers to communicate with the beamline system. The data acquisition computers need and have had the following services:

- Ability to monitor and record on tape beamline settings, status, intensity, etc.

- Ability to receive data from SWIC scanners including computed quantities and individual wire data.

- Ability to control devices.

- Ability to receive timing signals (e.g. beginning/end of spill) for synchronizing data acquisition hardware and software with the beam.

These functions are accomplished in the present system by means of (1) pulse trains; (2) 032 fifo module and (3) a variety of predets, timing fanouts, clock signals, etc. In many cases, the main reason for accumulating information from the control system is to have an accurate record of beamline conditions, settings, etc. However, some experiments, especially in the Neutrino areas, have transferred critical normalization information through the control system. Some of this information has been derived from special experiment-specific equipment. For these users, the beamline control system becomes an extension of their experiment's data acquisition system. This places demanding requirements on the beam line system's reliability, availability, and flexibility. The use of the beam line control system for *data acquisition* for experiments is discussed below.

Experience with these devices has been mixed. Nearly all experiments make use of pulse train and timing modules. However, only about 25% of the experiments have used the 032 module. This seems to be a result of the inconsistent performance of the hardware and the cumbersome nature of the information exchange protocol (again partly attributable to hardware limitations). Even successful users of the system complain about the large effort needed to support the hardware and software and are unhappy about the limited amount of data that can be transmitted. From the control system developer's point of view, the service consumes a lot of resources, contributing to system loading and puts a large maintenance load on the programming staff.

3.5.2 A Network Solution

Nearly all experimenters expressed the need for information interchange between their data acquisition computers and the beamline control system. The recent installation of DECnet connections among all the data acquisition VAXes has opened up a new approach to this problem. It will now be possible, as an alternative to the 032 connection, to establish a direct task-to-task communication between experiment VAXes and the beamline control system. This approach of supplying data through a computer network has the following implications:

- *The DECnet link in the experimental area must have high reliability and availability.* Major cables segments should be multiply connected via routers or bridges to take full advantage of the automatic failover and adaptive routing capabilities provided by DECnet. Any remaining 100Kb/s segment interconnections should be replaced with full bandwidth (optical?) links.

- All experiments who wish to use this service must have at least one computer that can run DECnet. (A few experiments have single PDP-11s that don't have enough memory to run DECnet along with their usual program load). The implications of this should be discussed with the Computing Department.

- A convenient software protocol must be designed to accomplish the data exchanges. On the experimenter's side, this protocol must conveniently connect to the typical VAX-FORTRAN environment.

- Services must be provided or developed to send data from the receiving VAX to the computer that writes the data tape if tape logging is required.

- If the system includes the ability to set devices from the experimenter's computer, complicated security issues must be resolved. Since the experimental computers are connected to DECnet and are frequently accessed by university collaborators from all over the world, the system will have to be somehow protected against unauthorized access. This same problem exists on the beamline computers themselves but the software environment is more controllable. The experimenter computers are managed by the experimenters themselves and practices with respect to security and access vary widely. (This problem is discussed again below).

It is the conclusion of the Task Force that this style of communication will resolve most of the difficulties of the present system and will win wide acceptance. We support the development of this system in conjunction with the experimenters and the Computing Department. The 032 link, however, must be retained until all users have successfully migrated to the DECnet link.

The committee members also believe that services of this kind will be more in demand in the future than they have been in the past. This will put a lot of pressure on the present CAMAC link so other solutions, such as special broadcast networks, in addition to DECnet ought to be considered to handle part of the load.

3.5.3 Volume of Data Transfer to Experiments

The amount of data transferred to an experiment varies widely from experiment to experiment and run to run. As an example of the possible needs of any future high demand experiment we can use Neutrino narrow band. An independent study of beamline needs was made for a future narrow band run based upon past experience and experimenter's needs. These experiments were making cross-section measurements where the incident flux was composed of neutral particles. As a requirement of the physics, almost all there was to know about the beam line was logged by the experimental computers after transfer from the beamline computer system. The magnet currents for the pre-target train and target train were monitored to fit flux measurements to the expected beam via TURTLE. Several primary and secondary flux measurement devices were recorded. Several beam position monitors, especially SWICs, were recorded. Finally, special devices were used, usually during special beam studies.

Modeling the possible future requirements on the types of flux monitoring done in the past and the existing dichromatic train, gives the estimated device reads of Table I. From this table, we see there are of order 20 reads for a spill plus of the order of 100 reads for each of several pings. For the last run, there were three pings but for the last dichromatic running there were five pings. Thus, there could be a future requirement of 300+ device reads per spill just for continuous flux measurement.

TABLE I

Representative Flux Monitoring Requirements

Magnet current readbacks	
4	pre-target train
6	target train
10	Total read for each ping
Collimator	
4	angle and momentum
1	hole
5	Total read each spill
Flux Monitoring	
4	pre-target
5	post-target in NC1
6	in NC2
6	downstream
21	Total read each ping
82	Total read each ping for each of 4 gates
TOTALS	
5	device reads each spill
92	device reads each ping

The beam position was monitored and recorded by means of split plate ion chambers, which are included above as flux monitoring devices, and by SWICs. Table II lists the possible SWIC requirements. These twelve possible SWICs would have one scan for each ping transferred to the experimental measurement.

TABLE II

Representative SWIC Requirements

4	pre-target train
4	post-target train
4	downstream

Special runs would also have to be conducted during the data taking to study the beam. These beam studies in the past included dump in runs, hole collimator runs and Cerenkov runs. The largest resource demands are made by the Cerenkov counter runs because the beam must be very stable and monitored to be so. Also some degree of control from the experiment of the beamline was required to control the Cerenkov counter at its far upstream location. Studies reported on elsewhere have shown that for a future Cerenkov counter to function in a higher energy secondary beam, that counter must be much more complex. This complexity translates into a far higher degree of needed control and much increased data readback per ping. One scheme has a CCD readout with 250k words readout and possibly transferred to the experimenters per ping. Other possible new beam monitoring hardware for better beam understanding could require devices with hundreds of words read out and transferred per ping, continually or for limited parts of the running.

In summary, we see that the amount of data coming from the beamline system to the experimenters for any future run comparable to a Neutrino narrow band run would involve the transfer of some 300 device reads plus the data from some 12 SWIC scanners at about 300 words each plus some large number of words from special purpose devices. While many of the devices are experiment specific, the community would be best served if they were controlled and read out via the beamline control system for three reasons. First, the devices are distributed within the beamline far from any experiment; second, they are often useful to the on-line beam monitoring system, and finally they are often collected for more than one experiment.

3.6 Connecting User Devices To The Control System

Increasingly there is available on the commercial market a growing variety of "smart" systems for purchase, or "smart" system building blocks. It is not a tenable position to argue that such systems should not be integrate-able into control systems at all! On the other hand, some such technologies, (e.g. GPIB) have been found by experience in the Accelerator Division to be extremely difficult to support. The Accelerator Division claims the reason for this is that the *so-called protocols such as GPIB are* ***not****, in fact, protocols at all. They are simply wire or connector specifications to first order*. It cannot be overemphasized that personnel hoping to integrate systems into an upgraded RD/Control system have a very heavy responsibility to communicate their plans and desires to a knowledgeable staff within the RD/Controls organization before heavy commitments are made. The RD/Controls organization must be able to support (at some level of arbitrariness) commercial "smart" systems and must begin to consider these problems soon. As mentioned earlier, it is not at all inappropriate for Laboratory management to request and require a uniformity of approach in these newer areas so that experience and developments in the Research and Accelerator Division are shared and not hidden.

Specifically, the Research Division should look at questions related to at least the following areas:

- GPIB devices (or, almost equivalently, intelligent devices using RS-232 such as the TI-550 family);

- Distributed intelligent processors/systems interfaced by DECnet/Ethernet/802.3 or 802.5;

- Systems developed by the Accelerator Division (GAS speakers).

3.7 System Security Issues

The security issue in a control system is primarily concerned with managing access to the individual devices which constitute the control environment and with preserving the integrity of critical files such as the master database. The type of access which is of concern is any operation upon a device (or a file) which changes its state; and, while this usually is a write operation, there exist devices which possess destructive read properties. A related, but less critical issue, is access to other system resources such as memory (both primary and secondary) and computer execution cycles.

In the recent past, protection against unauthorized access has addressed inadvertent rather than deliberate attempts at disruption; for example, a typing mistake in entering a device name. However, with the advent of networks and other connections into the system which reach outside of the laboratory boundaries, the possibility of a conscious attempt to bypass the access control mechanisms must be considered.

If a control system may be accessed by a widely distributed computer network, then unauthorized access to possibly critical devices becomes a definite concern. Such a "security" problem is an administrative question, because there is no simple compromise between a system which permits remote access for convenience and a "secure" system which does not, by design, permit remote access. A system which permits remote access can, with sufficient effort, always be compromised, no matter what mechanism is employed to provide "security".

The majority of the users wished to be able to access the control system for readings from any terminal within the laboratory and some, particularly the beamline physicists, wanted the same access from terminals at home. The most conservative group was the operators who wanted no access from outside the laboratory and were concerned about accesses from general laboratory terminals.

Only a few beamline physicists felt that access for control operations (which change the state of the system) from outside the laboratory was necessary but those that did expressed strong feelings on the issue, pointing out that the ability to change device settings from their homes would often save a trip in the middle of the night and would shorten the down period for a beamline problem. The alternative of having an operator change settings via a telephone connection was not satisfactory primarily because the lack of a display of readings, status, etc. made the process unnecessarily cumbersome. All of those who wished this class of remote access did feel that an explicit permit from an operator was an acceptable pre-condition.

Although primarily an implementation consideration, it would probably be difficult to allow control accesses from a terminal connected to an experimenter's computer while, at the same time, excluding the same type of access from an off-site terminal which was connected to the same computer via the network or the port selector.

There exist various means of managing access in a distributed computer system; but, inevitably, as the system security increases, the ease of access for legitimate users decreases. The administration must make a simple choice which states either that the system must remain uncompromised (this means NO remote access) or that remote access with the definite possibility of unauthorized access is tolerable.

3.8 Controls Hardware

The current Fermilab Experimental Areas control system is based upon a serial CAMAC link. This link was designed and installed prior to any formal serial link standard being published and as such suffers from a lack of compatibility with the now defined standard serial link. The serial CAMAC link is now over fifteen years old and has imposed several severe restrictions on the control system. One such restriction is the inability to generate LAM (look-at-me) signals for use in system interrupts. This lack of interrupts has forced the system to continually poll the system devices and thereby places a heavy load on the serial CAMAC link and on the control computers themselves. The present clock and data repeater system is outdated as well. Its high signal and power levels cause unwanted noise in experiment and other control system apparatus. Power down in one area causes all downstream areas to be nonfunctional. The Serial CAMAC Crate Controller as well as the clock and data repeating system should be replaced. The clock system should be upgraded in both its frequency and number of encoded times available. The beamline display consoles have also evolved over the years but the last major revision was done in 1976 with the installation of the 038 module. Terminal technology has come a long way since then with the advent of low cost color monitors and high resolution graphics. The current system does not have a suitable means to support this form of hardware and must do so if we are to take advantage of the currently available industry standards. The Terminals should be replaced with common and commercially available hardware. Terminal controllers will necessarily need to be replaced as well.

3.8.1 CAMAC Modules In General And MADCs

If we look at the list of CAMAC modules in use in the beamline system we find the last module to be designed and installed is now approaching four years old and the power supply controllers are now twelve years old. The lack of flexibility in the hardware has caused the inclusion of band-aid approaches to the various problems encountered in the beamline operations and, perhaps more seriously, has caused the addition of parallel systems to perform functions previously considered to be in the realm of the control system. These parallel systems must now be re-integrated into the system to provide the remote monitoring and control of these functions.

Systems such as power supplies which have a high impact on beam delivery in the experimental areas must be looked at both in regard to the control aspect and the readback of these devices. As mentioned above the control modules for these devices are now approaching their thirteenth year of service. Modules that were "state of the art" thirteen years ago are now only marginally acceptable in performance and in some instances are in fact unacceptable. The associated readback signals from power supplies attempts to send analog signals over hundreds of feet and have the control system resolve the signal at the millivolt level in electrically noisy environments. There are documented cases of current changes of a few amperes having caused intolerable beam motion at an experimenters apparatus. These changes amount to tens of millivolts at the MADC input! The same level of change from the controller to the power supply will again result in the unwanted beam motion. The present MADC analog readback system and D/A analog control system should be replaced. Two alternatives are possible and both should be considered relative to how their implementation impacts on total system requirements. The first alternative removes all long distance analog control and readback by putting analog signals at the power supply box only and controlling supplies with local intelligence in the power supply box remotely set up and monitored by digital means. The second alternative, although less attractive, replaces the present MADC system with up-to-date electronics and analog shielding methods.

We must begin to correct and up-date all of the beamline hardware modules if we are to have any hope of providing a system that matches the level of Physics being done at Fermilab. This must include the use of modules designed for the Accelerator where possible so we do not have to invest a large amount of time and effort re-designing and building similar hardware. The only hope of accomplishing this task is with a level of funding substantially higher than is currently provided.

3.8.2 Closed Loop Control

The concept of Closed Loop Control has taken on several different meanings in the realm of Fermilab control systems. One idea is that the setting of a device to a given position or current will automatically cause the control system to adjust the command signal until the readback satisfies the command. A second concept of closed loop control involves the use of beam position monitoring devices (SWICs, etc.)being used to control currents in steering magnets thereby keeping the beam centered on target. Both of these ideas are correct in their interpretation of closed loop control but each has a vastly different effect on the control system.

In the first instance if all the system elements were in perfect adjustment there would be no system intervention at all. A command would be matched by a perfect readback response. Since this is not the case the control computers would have to issue incremental changes until the command and readback matched. Since this would only happen when device settings were changed, it should represent a small overhead to the system. There are some obvious dangers in that if the device did not respond to the incremental change or if the increment were so large as to cause the readback to overshoot its desired point an unstable situation would occur and with it a continuous load on the system.

The second form of closed loop control has some variations as well. In some situations the controller predicts where the controlled device will be based upon previous settings and readbacks. It, in fact, "guesses" where the device will be and applies a corrective signal based upon that "guess". This is a predictive closed loop controller. The time period it is trying to correct for can be established as one of the closed loop parameters. This time period could be for the next beam cycle, the next second, the next millisecond, etc.. This is a discrete predictive closed loop controller. The other type of controller is a continuous closed loop controller. It is a classical linear system controller. The key is it operates in real time and could be thought of in terms of operational or linear amplifiers. This form of closed loop control would not be practical with any form of control system without having a dedicated high speed controller at the device hardware level. This device must have direct control over the device it is controlling without an intervening computer system to slow the system response.

Fermilab has applications for all types of closed loop control but the appropriate solution to each must be engineered into the systems. Simple set/readback closed loop control could be implemented at the power supply or hardware level. Algorithmic low repetition rate closed loop can be implemented at the control system level but it should serve to update parameters for device specific controllers so as not to pose a continuous load to the control system computers. Finally, real-time closed loop control can only be done with dedicated high speed hardware. The control system's role in this type of control should be purely monitoring and perhaps overall enabling or disabling of the closed loop function.

3.8.3 System Diagnostics

A system diagnostic module should be developed and installed in every controls crate. It should be capable of testing network as well as crate parameters. Along with this module, good system diagnostic software should be developed from the beginning using expert system software packages. Diagnostics should initially point to a failed crate or cable and eventually point to a malfunctioning module. Diagnostic messages should give precise information as to the cause of the failure.

3.8.4 Clock System

The Experimental Areas Clock system is used to synchronize the beamline related events to the Accelerator cycle. The original clock system consisted of a 1Mhz bipolar clock with phase reversal encoded event times. This system provided fifteen unique events which could be decoded by hardware to provide system timing information. Of the fifteen possible event times generally only five or six times were used. The changes brought about by Tevatron have forced the phase reversal clock rate to be reduced to 100Khz. This change was forced by the extended cycle time and the lack of depth in the hardware counter chains. The Accelerator Division has now gone to a different, more flexible, clock system. The new system employs a clock with 256 encoded events. (A block of 16 events on this new clock has already been reserved for use by the Research Division. More might be available.) The new clock would offer much greater flexibility in the control and timing of both beamline and experimental data acquisition events. We must begin to convert to the new style clock where ever possible and provide any new hardware with the ability to decode the new style clock. The Accelerator has also commissioned several other methods of broadcasting non-timing information. We should further investigate the possibility of using data provided on these other Accelerator broadcast systems in order to increase the reliability and ease of experimental data acquisition.

3.8.5 Other Concerns

While pointing to a few of the more noticeable problems above, one must by no means conclude that they are the only modules or systems in need of re-design. We must in fact decide on the general applicability of CAMAC as a data gathering system in the current state of electronics and standards. Although CAMAC is a reliable, mature standard, it should not be considered mandatory in future control system upgrades.

3.9 A Final Comment

Many of the requirements on the Research Division beamline control system are driven by the nature of the Fixed Target program and its users. The experimenter-user cannot be required or even expected to become an expert in the use of the control system. He must be supplied with a system which can satisfy the needs of his experiment (some of which may be unique) and which is easy to learn and convenient to use. It must provide him with all the beamline services required to run and monitor the beam, collect his data and aid in analyzing his experiment. It must be flexible enough to handle, at least at some level, new equipment he wants to connect to the system.

It has been observed by members of the Accelerator Division that the problems of "commissioning" new facilities (as opposed to operating them) place the greatest demand upon a control system. The Fixed Target program, with its great diversity, always has some commissioning activity going on and therefore requires a control system that has much greater capability than the present one.